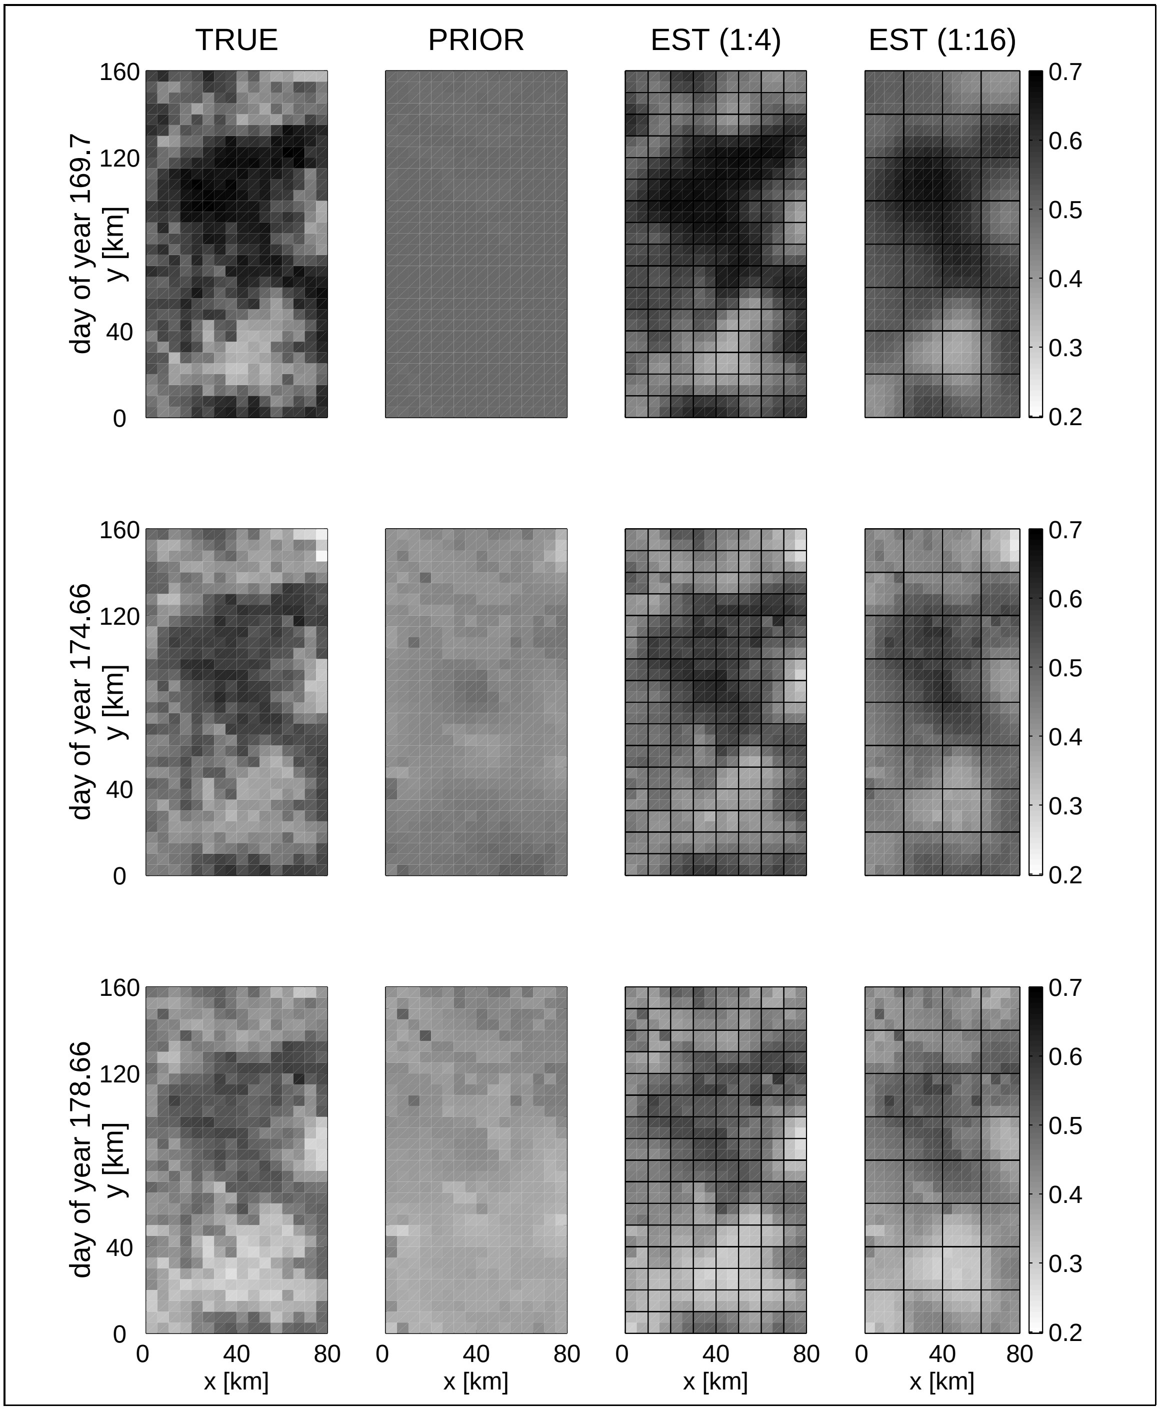

Figure 6.16: Top node saturation for the downscaling experiments. The first column shows the true top node saturation at three different times during the assimilation interval. In the second column, the prior top node saturation is depicted for the same three times. In columns three and four, the estimated top node saturation is shown for downscaling ratios of (1:4) and (1:16), resp. In the (1:4) scenario, each observation pixel covers four estimation pixels. The observation pixels are drawn with solid grid lines. The assimilation algorithm is clearly capable of capturing structures that are finer than the scale of the observations.

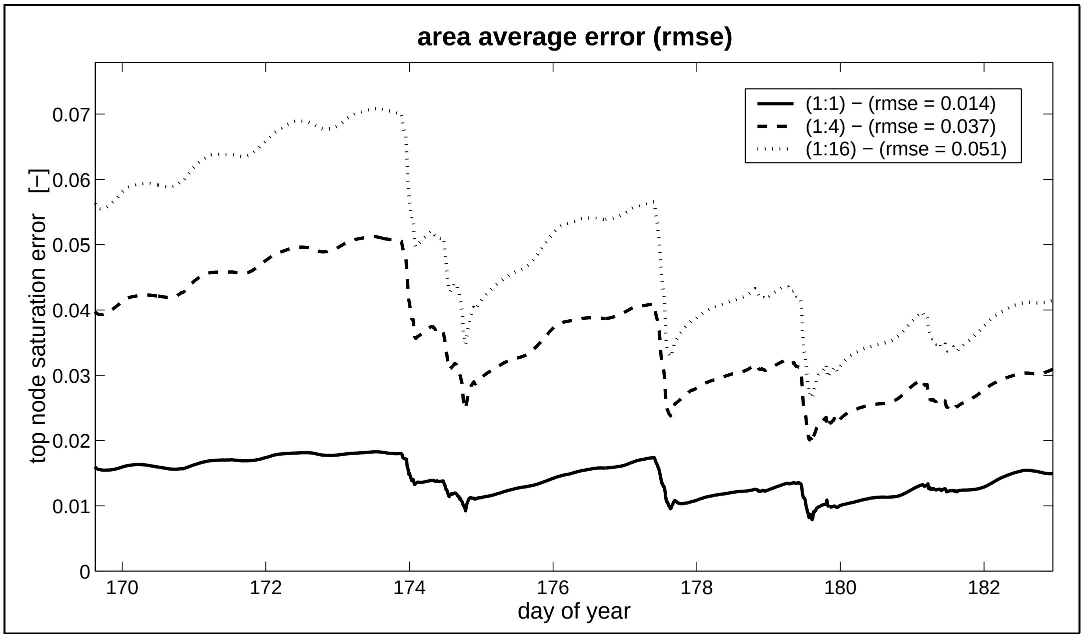

Figure 6.17: Area average errors for the downscaling experiments. The root-mean-square errors (rmse) in the estimated top node saturation with respect to the (synthetic) true fields are shown for Reference Experiment I (1:1) and two downscaling scenarios. The prior error is shown in Figure 6.4. In the legend we also indicate the temporal mean of the area average rmse. Note that the soil moisture errors are in terms of saturation (see Figure 6.4). As expected, the errors increase with decreasing resolution of the brightness data, but even for the coarser brightness images (1:16), the errors are still acceptable. The decrease with time in the errors of the two downscaling scenarios can be explained in the same way as the decrease in the prior error of Figure 6.4 (Section 6.1).

brightness images with resolutions of a few tens of kilometers are definitely useful even if the scale of interest is on the order of a few kilometers, provided we have fine-scale information on the micro-meteorologic forcings and on the model parameters.

Figure 6.17 offers a complementary view by showing time series of the area average top node saturation errors for all downscaling scenarios. The (1:1) downscaling scenario corresponds to Reference Experiment I of Section 6.1. Spatio-temporal averages of the errors are given in the legend. Note again that the errors are in terms of the saturation (see Figure 6.4). As expected, the area average error in the top node saturation increases with decreasing resolution of the brightness data. Even for the coarser (1:16) downscaling scenario, the root-mean-square error is always below 0.07, which translates into approximately 3% volumetric soil moisture. Of course the experiments so far have been assimilation exercises under ideal conditions, and the numbers are subject to change when field data are assimilated.

Finally, Figure 6.18 shows the reduced objective function for both downscaling scenarios. In both cases, the reduced objective function lies within about 1.5 standard deviations from the expected value. This confirms that the algorithm is working properly under the ideal

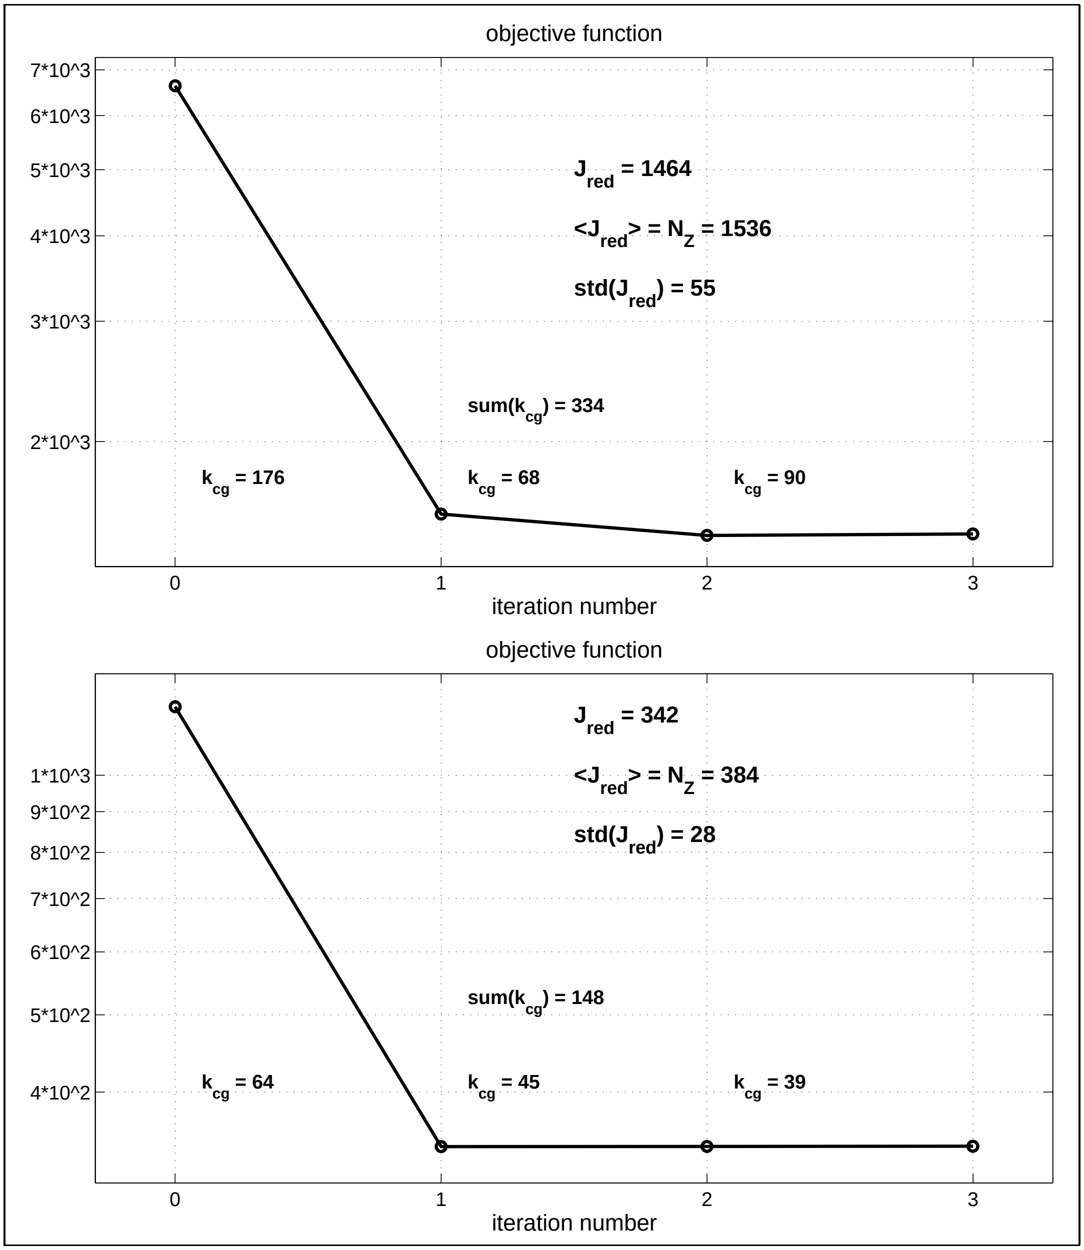

Figure 6.18: Objective function versus iteration number for the downscaling experiments. The upper and lower panels show the objective for the (1:4) and the (1:16) downscaling experiments, respectively. The reduced objective function after convergence is 1464 and 342, respectively. The number of data points is 1536 (384), which is also the expected value of the reduced objective function. The standard deviation of the reduced objective function is 55 (28). The values of k_{cg} indicate the number of linear combinations of representer functions that needed to be evaluated during the conjugate gradient iteration of the indirect representer approach (Chapter 8).

conditions of these experiments.

6.3.2 Determining the Horizontal Resolution

Ideally, we would like to determine the horizontal resolution of the model and hence of the estimation pixels only according to the availability of the model inputs such as the micro-meteorologic forcings or the soil and land cover parameters. In practice, however, the model resolution will also depend on the resolution of the brightness images, as well as on the computational feasibility. If we only get brightness data at $50km$ resolution, then it probably does not make much sense to work with a model resolution below $10km$. Clearly, a brightness average over $50km$ will have very little additional information when assimilated into a model of $1km$ resolution, and the assimilation procedure would effectively return the prior fields as the best estimate. Moreover, the computational effort increases with increasing resolution of the model (Section 8.2.2). Even if the model parameters are available at a very fine resolution, we may not be able to afford this fine a model resolution when we carry out the assimilation.

6.4 Satellite Repeat Cycle Experiments

6.4.1 Motivation and Experiment Design

Another question of practical relevance is the sensitivity of the estimates to the satellite revisit frequency. Since L-band microwave sensors require rather large antennae, and since the antenna size to achieve a given horizontal resolution increases with increasing orbit height, only polar orbiting satellites can be equipped with such sensors. This implies the disadvantage of a rather long revisit interval, which is typically on the order of three to four days for one polar orbiting satellite. By installing sensors on a few satellites, the repeat cycle may be somewhat reduced. However, with current technology and funding, a repeat cycle of one day, that is one image per day, appears to be the best one can hope for.

The most important factor determining soil moisture is certainly precipitation. But provided we have good observations of precipitation, most of the uncertainty in the soil moisture stems from poor knowledge of the initial condition and of the land surface fluxes. Now consider that soil moisture has a relatively long memory compared to soil temperature and other land surface variables. We can therefore hope to estimate the initial condition with rather infrequent observations, and therefore remove some of the uncertainty in the soil moisture estimates. The question is then how much the uncertainty in the land surface fluxes affects the soil moisture estimates. To examine this problem, we have conducted two experiments in which fewer brightness observations were assimilated.

In the first experiment with a 3-day brightness repeat cycle, only four (synthetic) brightness images at observation times 1, 4, 7, and 10 were assimilated. In another experiment, the assimilated data were further reduced to a 6-day repeat cycle, and only the measurements at observation times 1 and 7 were assimilated. See Figure 6.2 for details on the observation times. Otherwise the setup of the experiments is identical to the setup of Reference Experiment II. In particular, the repeat cycle sensitivity experiments were carried out with no downscaling and just one assimilation window covering the entire two-week period.

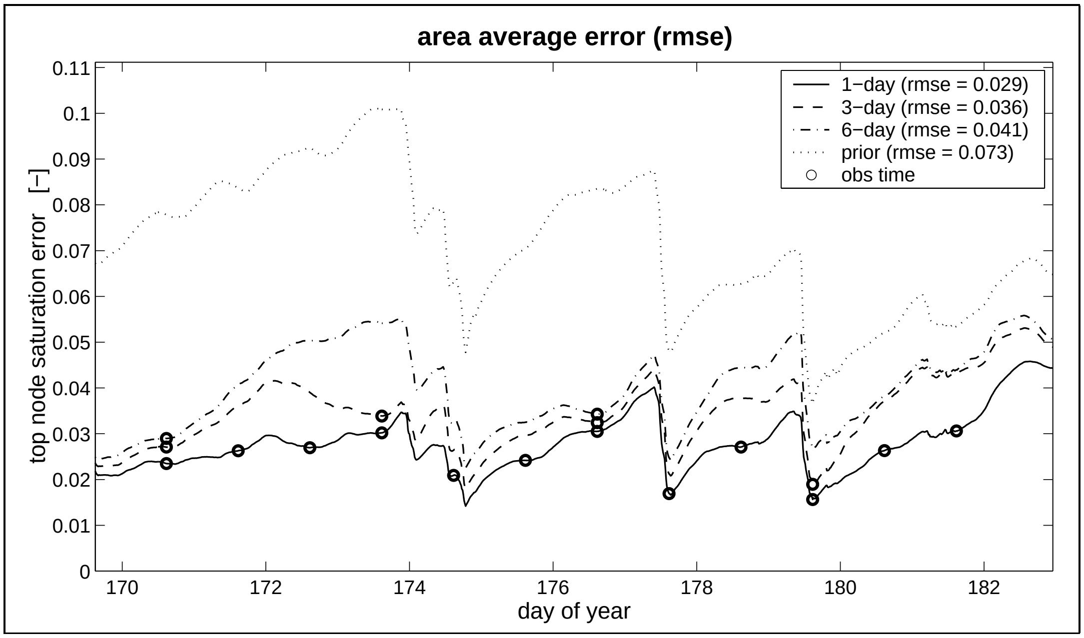

Figure 6.19: Area average errors for the repeat cycle experiments. The root-mean-square errors (rmse) of the estimated top node saturation with respect to the (synthetic) true fields are shown for Reference Experiment II (1-day repeat cycle) and two experiments with brightness repeat cycles of three and six days. The circles show the times at which observations are assimilated for each of the three experiments. In the legend we also indicate the temporal mean of the area average rmse. Note that the soil moisture errors are in terms of saturation (see Figure 6.4). As expected, the errors increase with decreasing availability of the brightness data.

6.4.2 Estimation of the True Fields

Figure 6.19 shows the area average top node saturation errors for the different brightness repeat cycles. Note again that the errors are in terms of saturation (see Figure 6.4). In the legend, temporal averages of the area average errors are also given. As expected, the time and area average error increases as we assimilate fewer brightness images.

At observation times 1 and 7, the errors for all experiments are nearly identical, because at these times brightness data are assimilated in all three experiments. Moreover, at observation times 4 and 10, the errors of the 3-day repeat cycle experiment and of Reference Experiment II are very close, because at these times brightness observations are assimilated in both experiments. After such shared observation times, however, the error of the 6-day and the 3-day repeat cycle experiments increases markedly. With observations available only every three or six days, we cannot expect to accurately estimate the model errors, which vary on the scale of their correlation time of 10 hours.

Since in all experiments we assimilate brightness data at the first observation time, we can reasonably well detect the initial condition for all repeat cycles. This explains why the estimates of all experiments are still much better than the prior trajectory. Even if we assimilate brightness data only every six days, we can reduce the prior error significantly, provided of course we have good precipitation data.

If we have only poor measurements of precipitation, the uncertainty in soil moisture will be dominated by the lack of precipitation information. Since the time scale of variability of soil moisture is governed by the interstorm periods, the minimum repeat cycle is then determined by the frequency of rainstorms. In the absence of accurate precipitation data, we must assimilate brightness observation at least once per interstorm period in order to get a satisfactory estimate of soil moisture. In Section 7.2, we examine whether it is feasible to estimate soil moisture even under the extreme condition of not having any quantitative precipitation data.

Finally, Figure 6.20 shows the reduced objective function for the repeat cycle experiments. In both cases, the value of the reduced objective function suggests that the algorithm works optimally.

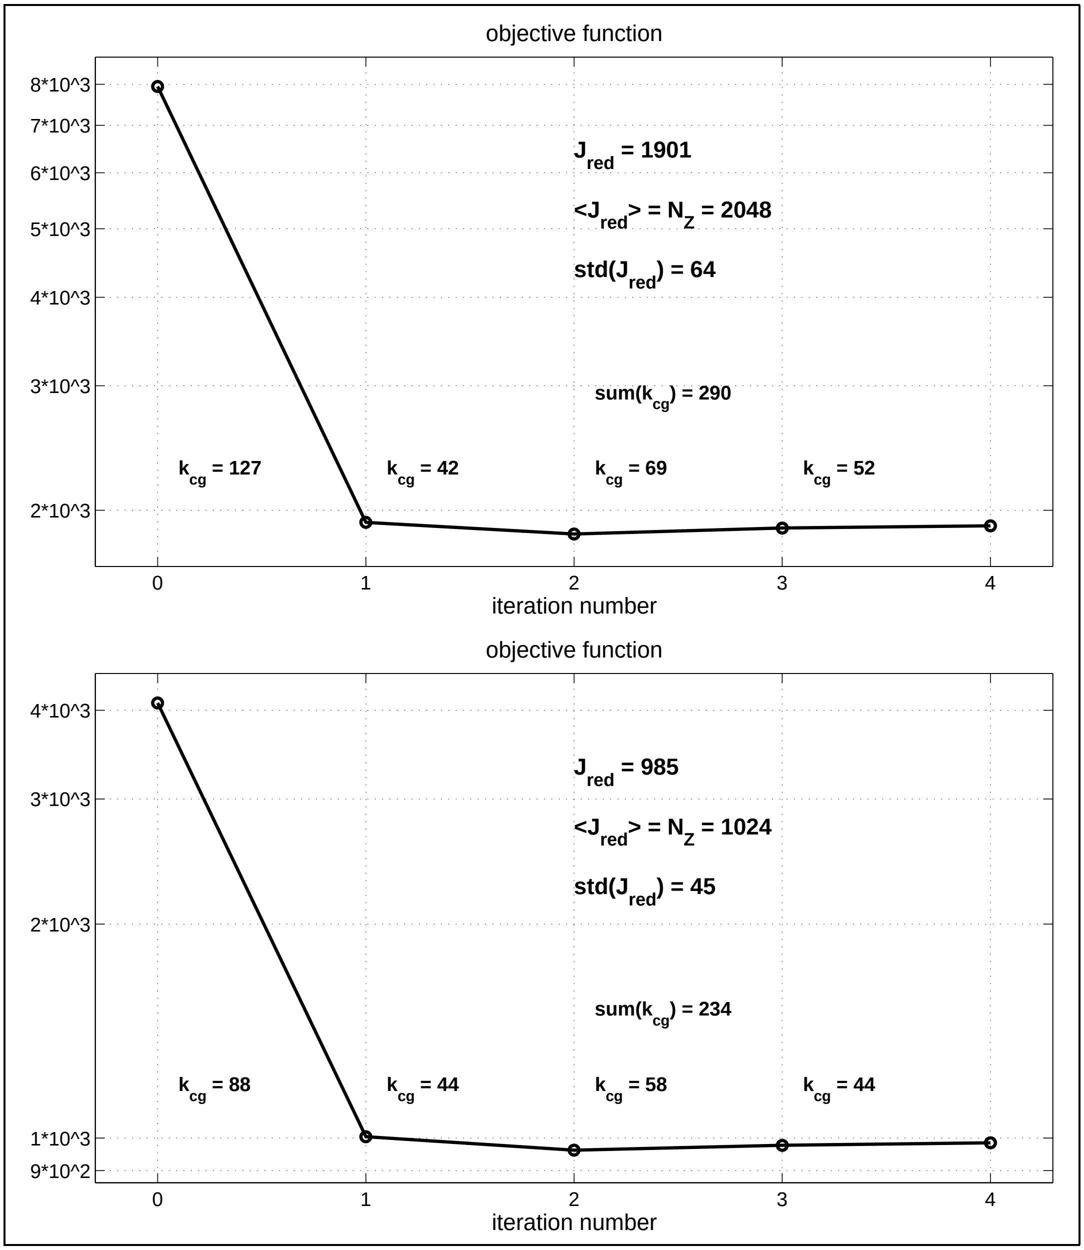

Figure 6.20: Objective function versus iteration number for the repeat cycle experiments. The upper and lower panels show the objective function for the 3-day and the 6-day brightness repeat cycle experiments, respectively. The reduced objective function after convergence is 1901 and 985, respectively. The number of data points is 2048 (1024), which is also the expected value of the reduced objective function. The standard deviation of the reduced objective function is 64 (45). The values of k_{cg} indicate the number of linear combinations of representer functions that needed to be evaluated during the conjugate gradient iteration of the indirect representer approach (Chapter 8).